

DATA ALLOCATION IN A DISTRIBUTED STORAGE SYSTEM**CROSS-REFERENCE TO RELATED APPLICATION**

This application is related to a U.S. patent application titled "Distributed Independent Cache
5 Memory," filed on even date, which is assigned to the assignee of the present application and which is incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates generally to data
10 storage, and specifically to data storage in distributed data storage entities.

BACKGROUND OF THE INVENTION

A distributed data storage system typically comprises cache memories that are coupled to a number of
15 disks wherein the data is permanently stored. The disks may be in the same general location, or be in completely different locations. Similarly, the caches may be localized or distributed. The storage system is normally used by one or more hosts external to the system.

20 Using more than one cache and more than one disk leads to a number of very practical advantages, such as protection against complete system failure if one of the caches or one of the disks malfunctions. Redundancy may be incorporated into a multiple cache or multiple disk
25 system, so that failure of a cache or a disk in the distributed storage system is not apparent to one of the external hosts, and has little effect on the functioning of the system.

While distribution of the storage elements has
30 undoubted advantages, the fact of the distribution typically leads to increased overhead compared to a local system having a single cache and a single disk. *Inter alia*, the increased overhead is required to manage the

increased number of system components, to equalize or attempt to equalize usage of the components, to maintain redundancy among the components, to operate a backup system in the case of a failure of one of the components, and to manage addition of components to, or removal of components from, the system. A reduction in the required overhead for a distributed storage system is desirable.

An article titled "Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web," by Karger et al., in the *Proceedings of the 29th ACM Symposium on Theory of Computing*, pages 654-663, (May 1997), whose disclosure is incorporated herein by reference, describes caching protocols for relieving "hot spots" in distributed networks. The article describes a hashing technique of consistent hashing, and the use of a consistent hashing function. Such a function allocates objects to devices so as to spread the objects evenly over the devices, so that there is a minimal redistribution of objects if there is a change in the devices, and so that the allocation is consistent, i.e., is reproducible. The article applies a consistent hashing function to read-only cache systems, i.e., systems where a client may only read data from the cache system, not write data to the system, in order to distribute input/output requests to the systems. A read-only cache system is used in much of the World Wide Web, where a typical user is only able to read from sites on the Web having such a system, not write to such sites.

An article titled "Differentiated Object Placement and Location for Self-Organizing Storage Clusters," by Tang et al., in *Technical Report 2002-32* of the University of California, Santa Barbara (November, 2002), whose disclosure is incorporated herein by reference, describes a protocol for managing a storage system where

components are added or removed from the system. The protocol uses a consistent hashing scheme for placement of small objects in the system. Large objects are placed in the system according to a usage-based policy.

5 An article titled "Compact, Adaptive Placement Schemes for Non-Uniform Capacities," by Brinkmann et al., in the August, 2002, *Proceedings of the 14th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, whose disclosure is incorporated herein by reference, describes
10 two strategies for distributing objects among a heterogeneous set of servers. Both strategies are based on hashing systems.

 U. S. patent 5,875,481 to Ashton, et al., whose disclosure is incorporated herein by reference, describes
15 a method for dynamic reconfiguration of data storage devices. The method assigns a selected number of the data storage devices as input devices and a selected number of the data storage devices as output devices in a predetermined input/output ratio, so as to improve data
20 transfer efficiency of the storage devices.

 U. S. patent 6,317,815 to Mayer, et al., whose disclosure is incorporated herein by reference, describes a method and apparatus for reformatting a main storage device of a computer system. The main storage device is
25 reformatted by making use of a secondary storage device on which is stored a copy of the data stored on the main device.

 U. S. patent 6,434,666 to Takahashi, et al., whose disclosure is incorporated herein by reference, describes
30 a memory control apparatus. The apparatus is interposed between a central processing unit (CPU) and a memory device that stores data. The apparatus has a plurality of cache memories to temporarily store data which is transferred between the CPU and the memory device, and a

cache memory control unit which selects the cache memory used to store the data being transferred.

U. S. patent 6,453,404 to Bereznyi, et al., whose disclosure is incorporated herein by reference, describes
5 a cache system that allocates memory for storage of data items by defining a series of small blocks that are uniform in size. The cache system, rather than an operating system, assigns one or more blocks for storage of a data item.

10

SUMMARY OF THE INVENTION

It is an object of some aspects of the present invention to provide a system for distributed data allocation.

5 In preferred embodiments of the present invention, a data distribution system comprises a plurality of data storage devices wherein data blocks may be stored. The data blocks are stored at logical addresses that are assigned to the data storage devices according to a
10 procedure which allocates the addresses among the devices in a manner that reduces the overhead incurred when a device is added to or removed from the system, and so as to provide a balanced access to the devices. The procedure typically distributes the addresses evenly
15 among the devices, regardless of the number of devices in the system. If a storage device is added to or removed from the system, the procedure reallocates the logical addresses between the new numbers of devices so that the balanced access is maintained. If a device has been
20 added, the procedure only transfers addresses to the added storage device. If a device has been removed, the procedure only transfers addresses from the removed storage device. In both cases, the only transfers of data that occur are of data blocks stored at the transferred
25 addresses. The procedure thus minimizes data transfer and associated management overhead when the number of storage devices is changed, or when the device configuration is changed, while maintaining the balanced access.

In some preferred embodiments of the present
30 invention, the procedure comprises a consistent hashing function. The function is used to allocate logical addresses for data block storage to the storage devices at initialization of the storage system. The same function is used to consistently reallocate the logical

addresses and data blocks stored therein when the number of devices in the system changes. Alternatively, the procedure comprises allocating the logical addresses between the devices according to a randomizing process at
5 initialization. The randomizing process generates a table giving a correspondence between specific logical addresses and the devices. The same randomizing process is used to reallocate the logical addresses and their stored data blocks on a change of storage devices

10 In some preferred embodiments of the present invention, the procedure comprises allocating two copies of a logical address to two separate storage devices, the two devices being used to store copies of a data block, so that the data block is protected against device
15 failure. The procedure spreads the data block copies uniformly across all the storage devices. On failure of any one of the devices, copies of data blocks of the failed device are still spread uniformly across the remaining devices, and are immediately available to the
20 system. Consequently, device failure has a minimal effect on the performance of the distribution system.

There is therefore provided, according to a preferred embodiment of the present invention, a method for data distribution, including:

25 distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices;

transferring the data to the storage devices in accordance with the logical addresses;

30 adding an additional storage device to the initial set, thus forming an extended set of the storage devices consisting of the initial set and the additional storage device; and

redistributing the logical addresses among the

storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and
5 without requiring a substantial transfer of the logical addresses among the storage devices in the initial set.

Preferably, redistributing the logical addresses consists of no transfer of the logical addresses between the storage devices in the initial set.

10 Preferably, distributing the logical addresses includes applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and redistributing the logical
15 addresses consists of applying the consistent hashing function to the extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

Alternatively, distributing the logical addresses
20 includes applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and redistributing the logical addresses consists of applying the randomizing function to the
25 extended set of storage devices so as to determine respective subsequent locations of the logical addresses among the extended set.

At least one of the storage devices preferably includes a fast access time memory; alternatively or
30 additionally, at least one of the storage devices preferably includes a slow access time mass storage device.

Preferably, the storage devices have substantially equal capacities, and distributing the logical addresses

includes distributing the logical addresses substantially evenly among the initial set, and redistributing the logical addresses consists of redistributing the logical addresses substantially evenly among the extended set.

5 Alternatively, a first storage device of the storage devices has a first capacity different from a second capacity of a second storage device of the storage devices, and distributing the logical addresses includes distributing the logical addresses substantially
10 according to a ratio of the first capacity to the second capacity, and redistributing the logical addresses includes redistributing the logical addresses substantially according to the ratio.

Preferably, distributing the logical addresses
15 includes allocating a specific logical address to a first storage device and to a second storage device, the first and second storage devices being different storage devices, and storing the data consists of storing a first copy of the data on the first storage device and a second
20 copy of the data on the second storage device.

The method preferably includes writing the data from a host external to the storage devices, and reading the data to the external host from the storage devices.

There is further provided, according to a preferred
25 embodiment of the present invention, an alternative method for distributing data, including:

distributing logical addresses among an initial set of storage devices so as provide a balanced access to the devices;

30 transferring the data to the storage devices in accordance with the logical addresses;

removing a surplus device from the initial set, thus forming a depleted set of the storage devices comprising the initial storage devices less the surplus storage

device; and

redistributing the logical addresses among the storage devices in the depleted set so as to cause logical addresses of the surplus device to be transferred
5 to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of logical addresses among the storage devices in the depleted set.

Preferably, redistributing the logical addresses
10 consists of no transfer of the logical addresses to the storage devices in the depleted set apart from the logical addresses of the surplus device.

Distributing the logical addresses preferably consists of applying a consistent hashing function to the
15 initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and redistributing the logical addresses preferably includes applying the consistent hashing function to the depleted set of storage devices
20 so as to determine respective subsequent locations of the logical addresses among the depleted set.

Alternatively, distributing the logical addresses consists of applying a randomizing function to the initial set of storage devices so as to determine
25 respective initial locations of the logical addresses among the initial set, and redistributing the logical addresses includes applying the randomizing function to the depleted set of storage devices so as to determine respective subsequent locations of the logical addresses
30 among the depleted set.

The storage devices preferably have substantially equal capacities, and distributing the logical addresses consists of distributing the logical addresses substantially evenly among the initial set, and

redistributing the logical addresses includes redistributing the logical addresses substantially evenly among the depleted set.

There is further provided, according to a preferred embodiment of the present invention, a method for distributing data among a set of storage devices, including:

applying a consistent hashing function to the set so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices;

forming subsets of the storage devices by subtracting the respective primary storage devices from the set;

applying the consistent hashing function to the subsets so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices; and

storing the data on the respective primary storage devices and a copy of the data on the respective secondary storage devices in accordance with the logical addresses.

There is further provided, according to a preferred embodiment of the present invention, a method for distributing data among a set of storage devices, including:

applying a randomizing function to the set so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices;

forming subsets of the storage devices by subtracting the respective primary storage devices from the set;

applying the randomizing function to the subsets so

as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices; and

storing the data on the respective primary storage devices and a copy of the data on the respective secondary storage devices in accordance with the logical addresses.

There is further provided, according to a preferred embodiment of the present invention, a data distribution system, including:

an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses; and

an additional storage device to the initial set, thus forming an extended set of the storage devices comprising the initial set and the additional storage device, the logical addresses being redistributed among the storage devices in the extended set so as to cause a portion of the logical addresses to be transferred from the storage devices in the initial set to the additional storage device, while maintaining the balanced access and without requiring a substantial transfer of the logical addresses among the storage devices in the initial set.

There is further provided, according to a preferred embodiment of the present invention, a data distribution system, including:

an initial set of storage devices among which are distributed logical addresses so as provide a balanced access to the devices, and wherein data is stored in accordance with the logical addresses; and

a depleted set of storage devices, formed by subtracting a surplus storage device from the initial set, the logical addresses being redistributed among the

storage devices in the depleted set so as to cause logical addresses of the surplus device to be transferred to the depleted set, while maintaining the balanced access and without requiring a substantial transfer of
5 the logical addresses among the storage devices in the depleted set.

Preferably, redistributing the logical addresses comprises no transfer of the logical addresses to the storage devices in the depleted set apart from the
10 logical addresses of the surplus device.

The distributed logical addresses are preferably determined by applying a consistent hashing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses
15 among the initial set, and redistributing the logical addresses preferably includes applying the consistent hashing function to the depleted set of storage devices so as to determine respective subsequent locations of the logical addresses among the depleted set.

Alternatively, the distributed logical addresses are determined by applying a randomizing function to the initial set of storage devices so as to determine respective initial locations of the logical addresses among the initial set, and redistributing the logical
20 addresses preferably includes applying the randomizing function to the depleted set of storage devices so as to determine respective subsequent locations of the logical addresses among the depleted set.

The storage devices preferably have substantially
30 equal capacities, and the distributed logical addresses are distributed substantially evenly among the initial set, and redistributing the logical addresses includes redistributing the logical addresses substantially evenly among the depleted set.

Alternatively or additionally, a first storage device included in the storage devices has a first capacity different from a second capacity of a second storage device included in the storage devices, and the distributed logical addresses are distributed substantially according to a ratio of the first capacity to the second capacity, and redistributing the logical addresses includes redistributing the logical addresses substantially according to the ratio.

Preferably, the distributed logical addresses include a specific logical address allocated to a first storage device and a second storage device, the first and second storage devices being different storage devices, and storing the data includes storing a first copy of the data on the first storage device and a second copy of the data on the second storage device.

The system preferably includes a memory having a table wherein is stored a correspondence between a plurality of logical addresses and a specific storage device in the initial set, wherein the plurality of logical addresses are related to each other by a mathematical relation.

There is further provided, according to a preferred embodiment of the present invention, a data distribution system, including:

a set of data storage devices to which is applied a consistent hashing function so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices; and

subsets of the storage devices formed by subtracting the respective primary storage devices from the set, the consistent hashing function being applied to the subsets so as to allocate the logical addresses to respective

secondary storage devices of the subsets while maintaining the balanced access to the devices, data being stored on the respective primary storage devices and a copy of the data being stored on the respective
5 secondary storage devices in accordance with the logical addresses.

There is further provided, according to a preferred embodiment of the present invention, a data distribution system, including:

10 a set of data storage devices to which is applied a randomizing function so as to allocate logical addresses to respective primary storage devices of the set and so as to provide a balanced access to the devices; and

subsets of the storage devices formed by subtracting
15 the respective primary storage devices from the set, the randomizing function being applied to the subsets so as to allocate the logical addresses to respective secondary storage devices of the subsets while maintaining the balanced access to the devices, data being stored on the
20 respective primary storage devices and a copy of the data being stored on the respective secondary storage devices in accordance with the logical addresses.

The present invention will be more fully understood from the following detailed description of the preferred
25 embodiments thereof, taken together with the drawings, a brief description of which is given below.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates distribution of data addresses among data storage devices, according to a preferred embodiment of the present invention;

5 Fig. 2 is a flowchart describing a procedure for allocating addresses to the devices of Fig. 1, according to a preferred embodiment of the present invention;

Fig. 3 is a flowchart describing an alternative procedure for allocating addresses to the devices of Fig. 1, according to a preferred embodiment of the present invention;

10 Fig. 4 is a schematic diagram illustrating reallocation of addresses when a storage device is removed from the devices of Fig. 1, according to a preferred embodiment of the present invention;

Fig. 5 is a schematic diagram illustrating reallocation of addresses when a storage device is added to the devices of Fig. 1, according to a preferred embodiment of the present invention;

20 Fig. 6 is a flowchart describing a procedure that is a modification of the procedure of Fig. 2, according to a preferred embodiment of the present invention;

Fig. 7 is a schematic diagram which illustrates a fully mirrored distribution of data for the devices of Fig. 1, according to a preferred embodiment of the present invention; and

Fig. 8 is a flowchart describing a procedure for performing the distribution of Fig. 7, according to a preferred embodiments of the present invention.

30

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Reference is now made to Fig. 1, which illustrates distribution of data addresses among data storage devices, according to a preferred embodiment of the present invention. A storage system 12 comprises a plurality of separate storage devices 14, 16, 18, 20, and 22, also respectively referred to herein as storage devices B_1 , B_2 , B_3 , B_4 , and B_5 , and collectively as devices B_n . It will be understood that system 12 may comprise substantially any number of physically separate devices, and that the five devices B_n used herein are by way of example. Devices B_n comprise any components wherein data 34, also herein termed data D, may be stored, processed, and/or serviced. Examples of devices B_n comprise random access memory (RAM) which has a fast access time and which are typically used as caches, disks which typically have a slow access time, or any combination of such components. A host 24 communicates with system 12 in order to read data from, or write data to, the system. A central processing unit (CPU) 26, using a memory 28, manages system 12, and allocates data D to devices B_n . The allocation of data D by CPU 26 to devices B_n is described in more detail below.

Data D is processed in devices B_n at logical block addresses (LBAs) of the devices by being written to the devices from host 24 and/or read from the devices by host 24. At initialization of system 12 CPU 26 distributes the LBAs of devices B_n among the devices using one of the pre-defined procedures described below. CPU 26 may then store data D at the LBAs.

In the description of the procedures hereinbelow, devices B_n are assumed to have substantially equal capacities, where the capacity of a specific device is a function of the device type. For example, for devices

that comprise mass data storage devices having slow access times, such as disks, the capacity is typically defined in terms of quantity of data the device may store. For devices that comprise fast access time
5 memories, such as are used in caches, the capacity is typically defined in terms of throughput of the device. Those skilled in the art will be able to adapt the procedures when devices B_n have different capacities, in which case ratios of the capacities are typically used to
10 determine the allocations. The procedures allocate the logical stripes to devices B_n so that balanced access to the devices is maintained, where balanced access assumes that taken over approximately $10,000 \times N$ transactions with devices B_n , the fraction of capacities of devices B_n used
15 are equal to within approximately 1%, where N is the number of devices B_n , the values being based on a Bernoulli distribution.

Fig. 2 is a flowchart describing a procedure 50 for allocating LBAs to devices B_n , according to a preferred
20 embodiment of the present invention. The LBAs are assumed to be grouped into k logical stripes/tracks, hereinbelow termed stripes 36 (Fig. 1), which are numbered 1, ..., k , where k is a whole number. Each logical stripe comprises one or more consecutive LBAs, and all the stripes have
25 the same length. Procedure 50 uses a randomizing function to allocate a stripe s to devices B_n in system 12. The allocations determined by procedure 50 are stored in a table 32 of memory 28.

In an initial step 52, CPU 26 determines an initial
30 value of s , the total number T_d of active devices B_n in system 12, and assigns each device B_n a unique integral identity between 1 and T_d . In a second step 54, the CPU generates a random integer R between 1 and T_d , and allocates stripe s to the device B_n corresponding to R .

In a third step 56, the allocation determined in step 54 is stored in table 32. Procedure 50 continues, in a step 58, by incrementing the value of s , until all stripes of devices B_n have been allocated, i.e., until $s > k$, at which point procedure 50 terminates.

Table I below is an example of an allocation table generated by procedure 50, for system 12, wherein $T_d = 5$. The identifying integers for each device B_n , as determined by CPU 26 in step 52, are assumed to be 1 for B_1 , 2 for B_2 , ... , 5 for B_5 .

Stripe s	Random Number R	Device B_s
1	3	B_3
2	5	B_5
...
6058	2	B_2
6059	2	B_2
6060	4	B_4
6061	5	B_5
6062	3	B_3
6063	5	B_5
6064	1	B_1
6065	3	B_3
6066	2	B_2
6067	3	B_3
6068	1	B_1
6069	2	B_2
6070	4	B_4
6071	5	B_5
6072	4	B_4
6073	1	B_1
6074	5	B_5
6075	3	B_3

6076	1	B_1
6077	2	B_2
6078	4	B_4
...

Table I

Fig. 3 is a flowchart showing steps of a procedure 70 using a consistent hashing function to allocate stripes to devices B_n , according to an alternative preferred embodiment of the present invention. In an initial step 72, CPU 26 determines a maximum number N of devices B_n for system 12, and a number of points k for each device. The CPU then determines an integer M , such that $M \gg N \cdot k$.

In a second step 74, CPU 26 determines N sets J_n of k random values S_{ab} , each set corresponding to a possible device B_n , as given by equations (1):

$$\begin{aligned}
 J_1 &= \{S_{11}, S_{12}, \dots, S_{1k}\} \text{ for device } B_1; \\
 J_2 &= \{S_{21}, S_{22}, \dots, S_{2k}\} \text{ for device } B_2; \\
 &\dots \\
 J_N &= \{S_{N1}, S_{N2}, \dots, S_{Nk}\} \text{ for device } B_N.
 \end{aligned} \tag{1}$$

Each random value S_{ab} is chosen from $\{0, 1, 2, \dots, M-1\}$, and the value of each S_{ab} may not repeat, i.e., each value may only appear once in all the sets. The sets of random values are stored in memory 28.

In a third step 76, for each stripe s CPU 26 determines a value of $s \bmod(M)$ and then a value of $F(s \bmod(M))$, where F is a permutation function that reassigns the value of $s \bmod(M)$ so that in a final step 78 consecutive stripes will generally be mapped to different devices B_n .

In final step 78, the CPU finds, typically using an

iterative search process, the random value chosen in step 74 that is closest to $F(s \bmod(M))$. CPU 26 then assigns the device B_n of the random value to stripe s , according to equations (1).

5 It will be appreciated that procedure 70 illustrates one type of consistent hashing function, and that other such functions may be used by system 12 to allocate LBAs to devices operating in the system. All such consistent hashing functions are assumed to be comprised within the
10 scope of the present invention.

Procedure 70 may be incorporated into memory 28 of system 12 (Fig. 1), and the procedure operated by CPU 26 when allocation of stripes s are required, such as when data is to be read from or written to system 12.
15 Alternatively, a table 30 of the results of applying procedure 70, generally similar to the first and last columns of Table I, may be stored in memory 28, and accessed by CPU 26 as required.

Fig. 4 is a schematic diagram illustrating
20 reallocation of stripes when a storage device is removed from storage system 12, according to a preferred embodiment of the present invention. By way of example, device B_3 is assumed to be no longer active in system 12 at a time $t=1$, after initialization time $t=0$, and the
25 stripes initially allocated to the device, and any data stored therein, are reallocated to the depleted set of devices B_1, B_2, B_4, B_5 of the system. Device B_3 may be no longer active for a number of reasons known in the art, such as device failure, or the device becoming surplus to
30 the system, and such a device is herein termed a surplus device. The reallocation is performed using procedure 50 or procedure 70, preferably according to the procedure that was used at time $t=0$. As is illustrated in Fig. 4, and as is described below, stripes from device B_3 are

substantially evenly redistributed among devices B_1 , B_2 , B_4 , B_5 .

If procedure 50 (Fig. 2) is applied at $t=1$, the procedure is applied to the stripes of device B_3 , so as to randomly assign the stripes to the remaining active devices of system 12. In this case, at step 52 the total number of active devices $T_d = 4$, and identifying integers for each active device B_n are assumed to be 1 for B_1 , 2 for B_2 , 4 for B_4 , 3 for B_5 . CPU 26 generates a new table, corresponding to the first and last columns of Table II below for the stripes that were allocated to B_3 at $t=0$, and the stripes are reassigned according to the new table. Table II illustrates reallocation of stripes for device B_3 (from the allocation shown in Table I).

15

Stripe s	Device B_s $t=0$	Random Number R $t=1$	Device B_s $t=1$
1	B_3	1	B_1
2	B_5		B_5
...
6058	B_2		B_2
6059	B_2		B_2
6060	B_4		B_4
6061	B_5		B_5
6062	B_3	3	B_5
6063	B_5		B_5
6064	B_1		B_1
6065	B_3	2	B_2
6066	B_2		B_2
6067	B_3	3	B_5
6068	B_1		B_1
6069	B_2		B_2
6070	B_4		B_4

6071	B ₅		B ₅
6072	B ₄		B ₄
6073	B ₁		B ₁
6074	B ₅		B ₅
6075	B ₃	4	B ₄
6076	B ₁		B ₁
6077	B ₂		B ₂
6078	B ₄		B ₄
...

Table II

It will be appreciated that procedure 50 only generates transfer of stripes from the device that is no longer active in system 12, and that the procedure reallocates the stripes, and any data stored therein, substantially evenly over the remaining active devices of the system. No reallocation of stripes occurs in system 12 other than stripes that were initially allocated to the device that is no longer active. Similarly, no transfer of data occurs other than data that was initially in the device that is no longer active. Also, any such transfer of data may be performed by CPU 26 transferring the data directly from the inactive device to the reallocated device, with no intermediate device needing to be used.

Similarly, by consideration of procedure 70 (Fig. 3), it will be appreciated that procedure 70 only generates transfer of stripes, and reallocation of data stored therein, from the device that is no longer active in system 12, i.e., device B₃. Procedure 70 reallocates the stripes (and thus their data) from B₃ substantially evenly over the remaining devices B₁, B₂, B₄, B₅ of the system, no reallocation of stripes or data occurs in system 12 other than stripes/data that were initially in

B₃, and such data transfer as may be necessary may be performed by direct transfer to the remaining active devices. It will also be understood that if B₃ is returned to system 12 at some future time, the allocation
5 of stripes after procedure 70 is implemented is the same as the initial allocation generated by the procedure.

Fig. 5 is a schematic diagram illustrating reallocation of stripes when a storage device is added to storage system 12, according to a preferred embodiment of
10 the present invention. By way of example, a device 23, also herein termed device B₆, is assumed to be active in system 12 at time t=2, after initialization time t=0, and some of the stripes initially allocated to an initial set of devices B₁, B₂, B₃, B₄, B₅, and any data stored
15 therein, are reallocated to device B₆. The reallocation is performed using procedure 70 or a modification of procedure 50 (described in more detail below with reference to Fig. 6), preferably according to the procedure that was used at time t=0. As is illustrated in
20 Fig. 5, and as is described below, stripes from devices B₁, B₂, B₃, B₄, B₅ are substantially evenly removed from the devices and are transferred to device B₆. B₁, B₂, B₃, B₄, B₅, B₆ act as an extended set of the initial set.

Fig. 6 is a flowchart describing a procedure 90 that
25 is a modification of procedure 50 (Fig. 2), according to an alternative preferred embodiment of the present invention. Apart from the differences described below, procedure 90 is generally similar to procedure 50, so that steps indicated by the same reference numerals in
30 both procedures are generally identical in implementation. As in procedure 50, procedure 90 uses a randomizing function to allocate stripes s to devices B_n in system 12, when a device is added to the system. The allocations determined by procedure 90 are stored in

table 32 of memory 28.

Assuming procedure 50 is applied at $t=2$, at step 52 the total number of active devices $T_d = 6$, and identifying integers for each active device B_n are assumed to be 1 for B_1 , 2 for B_2 , 3 for B_3 , 4 for B_4 , 5 for B_5 , 6 for B_6 . In a step 91 CPU 26 determines a random integer between 1 and 6.

In a step 92, the CPU determines if the random number corresponds to one of the devices present at time $t=0$. If it does correspond, then CPU 26 returns to the beginning of procedure 90 by incrementing stripe s , via step 58, and no reallocation of stripe s is made. If it does not correspond, i.e., the random number is 6, corresponding to device B_6 , the stripe is reallocated to device B_6 . In step 56, the reallocated location is stored in table 32. Procedure 90 then continues to step 58. Table III below illustrates the results of applying procedure 90 to the allocation of stripes given in Table II.

Stripe s	Device B_s $t=0$	Random Number R $t=2$	Device B_s $t=2$
1	B_3	6	B_6
2	B_5	4	B_5
...
6058	B_2	5	B_2
6059	B_2	3	B_2
6060	B_4	5	B_4
6061	B_5	6	B_6
6062	B_3	3	B_5
6063	B_5	1	B_5
6064	B_1	3	B_1
6065	B_3	1	B_2

6066	B ₂	6	B ₆
6067	B ₃	4	B ₅
6068	B ₁	5	B ₁
6069	B ₂	2	B ₂
6070	B ₄	1	B ₄
6071	B ₅	5	B ₅
6072	B ₄	2	B ₄
6073	B ₁	4	B ₁
6074	B ₅	5	B ₅
6075	B ₃	1	B ₄
6076	B ₁	3	B ₁
6077	B ₂	6	B ₆
6078	B ₄	1	B ₄
...

Table III

It will be appreciated that procedure 90 only generates transfer of stripes, and thus reallocation of data, to device B₆. The procedure reallocates the stripes to B₆ by transferring stripes, substantially evenly, from devices B₁, B₂, B₃, B₄, B₅ of the system, and no transfer of stripes, or data stored therein, occurs in system 12 other than stripes/data transferred to B₆. Any such data transfer may be made directly to device B₆, without use of an intermediate device B_n.

It will also be appreciated that procedure 70 may be applied when device B₆ is added to system 12. Consideration of procedure 70 shows that similar results to those of procedure 90 apply, i.e., that there is only reallocation of stripes, and data stored therein, to device B₆. As for procedure 90, procedure 70 generates substantially even reallocation of stripes/data from the other devices of the system.

Fig. 7 is a schematic diagram which illustrates a

fully mirrored distribution of data D in storage system 12 (Fig. 1), and Fig. 8 is a flowchart illustrating a procedure 100 for performing the distribution, according to preferred embodiments of the present invention.

5 Procedure 100 allocates each specific stripe to a primary device B_{n1} , and a copy of the specific stripe to a secondary device B_{n2} , $n1 \neq n2$, so that each stripe is mirrored. To implement the mirrored distribution, in a first step 102 of procedure 100, CPU 26 determines

10 primary device B_{n1} for locating a stripe using procedure 50 or procedure 70. In a second step 104, CPU 26 determines secondary device B_{n2} for the stripe using procedure 50 or procedure 70, assuming that device B_{n1} is not available. In a third step 106, CPU 26 allocates

15 copies of the stripe to devices B_{n1} and B_{n2} , and writes the device identities to a table 34 in memory 28, for future reference. CPU 26 implements procedure 100 for all stripes 36 in devices B_n .

Table IV below illustrates devices B_{n1} and B_{n2}

20 determined for stripes 6058 - 6078 of Table I, where steps 102 and 104 use procedure 50.

Stripe	Device B_{n1}	Device B_{n2}
6058	B_2	B_4
6059	B_2	B_5
6060	B_4	B_2
6061	B_5	B_4
6062	B_3	B_1
6063	B_5	B_4
6064	B_1	B_3
6065	B_3	B_4
6066	B_2	B_5
6067	B_3	B_1
6068	B_1	B_3

6069	B ₂	B ₅
6070	B ₄	B ₁
6071	B ₅	B ₃
6072	B ₄	B ₂
6073	B ₁	B ₃
6074	B ₅	B ₁
6075	B ₃	B ₅
6076	B ₁	B ₃
6077	B ₂	B ₄
6078	B ₄	B ₁

Table IV

If any specific device B_n becomes unavailable, so that only one copy of the stripes on the device is available in system 12, CPU 26 may implement a procedure similar to procedure 100 to generate a new second copy of the stripes that were on the unavailable device. For example, if after allocating stripes 6058 - 6078 according to Table IV, device B₃ becomes unavailable, copies of stripes 6062, 6065, 6067, and 6075, need to be allocated to new devices in system 12 to maintain full mirroring. Procedure 100 may be modified to find the new device of each stripe by assuming that the remaining device, as well as device B₃, is unavailable. Thus, for stripe 6062, CPU 26 assumes that devices B₁ and B₃ are unavailable, and determines that instead of device B₃ the stripe should be written to device B₄. Table V below shows the devices that the modified procedure 100 determines for stripes 6058, 6060, 6062, 6065, 6072, and 6078, when B₃ becomes unavailable.

Stripe s	Device B _{n1}	Device B _{n2}
6062	B ₁	B ₂
6065	B ₄	B ₅

6067	B ₁	B ₄
6075	B ₅	B ₂

Table V

It will be appreciated that procedure 100 spreads locations for stripes 36 substantially evenly across all devices B_n, while ensuring that each pair of copies of any particular stripe are on different devices, as is illustrated in Fig. 7. Furthermore, the even distribution of locations is maintained even when one of devices B_n, becomes unavailable. Either copy, or both copies, of any particular stripe may be used when host 24 communicates with system 12. It will also be appreciated that in the event of one of devices B_n becoming unavailable, procedure 100 regenerates secondary locations for copies of stripes 36 that are evenly distributed over devices B_n.

Referring back to Fig. 1, it will be understood that the sizes of tables 30, 32, or 34 are a function of the number of stripes in system 12, as well as the number of storage devices in the system. Some preferred embodiments of the present invention reduce the sizes of tables 30, 32, or 34 by duplicating some of the entries of the tables, by relating different stripes mathematically. For example, if system 12 comprises 2,000,000 stripes, the same distribution may apply to every 500,000 stripes, as illustrated in Table VI below. Table VI is derived from Table I.

Stripe s	Stripe s	Stripe s	Stripe s	Device B _s
1	500,001	1,000,001	1,500,001	B ₃
2	500,002	1,000,002	1,500,002	B ₅
...
6059	506,059	1,006,059	1,506,059	B ₂

6060	506,060	1,006,060	1,506,060	B ₄
...

Table VI

It will be appreciated that procedures such as those described above may be applied substantially
5 independently to different storage devices, or types of devices, of a storage system. For example, a storage system may comprise a distributed fast access cache coupled to a distributed slow access mass storage. Such a storage system is described in more detail in the U. S.
10 Application titled "Distributed Independent Cache Memory," filed on even date, and assigned to the assignee of the present invention. The fast access cache may be assigned addresses according to procedure 50 or modifications of procedure 50, while the slow access mass
15 storage may be assigned addresses according to procedure 70 or modifications of procedure 70.

It will thus be appreciated that the preferred embodiments described above are cited by way of example, and that the present invention is not limited to what has
20 been particularly shown and described hereinabove. Rather, the scope of the present invention includes both combinations and subcombinations of the various features described hereinabove, as well as variations and modifications thereof which would occur to persons
25 skilled in the art upon reading the foregoing description and which are not disclosed in the prior art.